Minimum Rates of Approximate Sufficient Statistics

Masahito Hayashi[†], *Fellow, IEEE*, and Vincent Y. F. Tan[‡], *Senior Member, IEEE*

arXiv:1612.02542v1 [cs.IT] 8 Dec 2016

Abstract

Given a sufficient statistic for a parametric family of distributions, one can estimate the parameter without access to the data itself. However, the memory or code size for storing the sufficient statistic may nonetheless still be prohibitive. Indeed, for n independent data samples drawn from a k-nomial distribution with $d = k-1$ degrees of freedom, the length of the code scales as $d \log n + O(1)$. In many applications though, we may not have a useful notion of sufficient statistics (e.g., when the parametric family is not an exponential family) and also may not need to reconstruct the generating distribution exactly. By adopting a Shannon-theoretic approach in which we consider allow a small error in estimating the generating distribution, we construct various notions of *approximate sufficient statistics* and show that the code length can be reduced to $\frac{d}{2} \log n + O(1)$. We also note that the locality assumption that is used to describe the notion of local approximate sufficient statistics when the parametric family is not an exponential family can be dispensed of. We consider errors measured according to the relative entropy and variational distance criteria. For the code construction parts, we leverage Rissanen's minimum description length principle, which yields a non-vanishing error measured using the relative entropy. For the converse parts, we use Clarke and Barron's asymptotic expansion for the relative entropy of a parametrized distribution and the corresponding mixture distribution. The limitation of this method is that only a weak converse for the variational distance can be shown. We develop new techniques to achieve vanishing errors and we also prove strong converses for all our statements. The latter means that even if the code is allowed to have a non-vanishing error, its length must still be at least $\frac{d}{2} \log n$.

Index Terms

Approximate sufficient statistics, Minimum rates, Memory size reduction, Minimum description length, Exponential families, Pythagorean theorem, Strong converse

I. Introduction

The notion of *sufficient statistics* is a fundamental and ubiquitous concept in statistics and information theory [1], [2]. Consider a random variable $X \in \mathcal{X}$ whose distribution $P_{X|Z=z}$ depends on an unknown parameter $z \in \mathcal{Z}$. Typically in detection and estimation problems, we are interested in learning the unknown parameter z. In this case, it is often unnecessary to use the full dataset X for this purpose. Rather a function of the data $Y = f(X) \in \mathcal{Y}$ usually suffices. If there is no loss in the performance of learning Z given Y relative to the case when one is given X, then Y is called a sufficient statistic relative to the family $\{P_{X|Z=z}\}_{z\in\mathcal{Z}}$. We may then write

$$P_{X|Z=z}(x) = \sum_{y\in\mathcal{Y}} P_{X|Y}(x|y) P_{Y|Z=z}(y), \quad \forall\, (x,z) \in \mathcal{X} \times \mathcal{Z} \tag{1}$$

or more simply that $X - Y - Z$ forms a Markov chain in this order. Because Y is a function of X, it is also true that $I(Z;X) = I(Z;Y)$ (where $I(Z;X)$ is the mutual information [2] between Z and X). This intuitively means that the sufficient statistic Y provides a much information about the parameter Z as the original data X does.

For concreteness in our discussions, we often (but not always) regard the family $\{P_{X|Z=z}\}_{z\in\mathcal{Z}}$ as an exponential family [3], i.e., $P_{X|Z=z} \propto \exp\left(\sum_i z_i Y_i(x)\right)$. This class of distributions is parametrized by a set of natural parameters $\{z_i\}$ and a set of natural statistics $\{Y_i(x)\}$, which is a function of the data. The natural statistics or maximum likelihood estimator (MLE) are known to be sufficient statistics of the exponential family. In many applications, large datasets are prevalent. In particular, the one-shot model described above will be replaced by an n-shot one in which the dataset consists of n independent and identically distributed (i.i.d.) random variables $X^n = (X_1, X_2, \ldots, X_n)$ each distributed according to $P_{X|Z=z}$ where the exact z is unknown. If the support of X is finite, the distribution can be regarded as a k-nomial distribution (a discrete distribution taking on at most k values) and the the empirical distribution or type [4] of X^n is a sufficient statistic for learning z. However, as is well known [4], the number of types with denominator n on an alphabet with k values is $\binom{n+k-1}{k-1} = \Theta(n^{k-1})$. Thus, the memory size required to learn the parameter z or distribution $P_{X|Z=z}$ is at least $\Theta(n^{k-1})$ if the (index of the) type is stored. The exponent $d = k-1$ here is the number of degrees of freedom in the distribution family, i.e., the dimensionality of the space $\mathcal{Z}$ that the parameter z belongs to. Can we do better than a memory size of $\Theta(n^d)$? The answer to this question depends on the strictness of the recoverability condition of $P_{X|Z=z}$. If $P_{X|Z=z}$ is to be recovered *exactly*, then the Markov condition in (1) is necessary and no reduction of the memory size is possible. However, if $P_{X|Z=z}$ is to be recovered only *approximately*, we can indeed reduce the rate of the memory size. This is one motivation for the current paper.

[†]Masahito Hayashi is with the Graduate School of Mathematics, Nagoya University, and the Center for Quantum Technologies (CQT), National University of Singapore (Email: masahito@math.nagoya-u.ac.jp).

[‡]Vincent Y. F. Tan is with the Department of Electrical and Computer Engineering and the Department of Mathematics, National University of Singapore (Email: vtan@nus.edu.sg).

In addition, going beyond exponential families, for general distribution families, we do not necessarily have a useful and universal notion of sufficient statistics. Thus, we often focus on *local asymptotic sufficient statistics* by relaxing the condition for sufficient statistics. For example, under suitable regularity conditions [5]–[7], the MLE forms a set of local asymptotic sufficient statistics. However, to the best of our knowledge, there is no prior work that discusses the required memory size is required for asymptotic sufficient statistics. To address this issue, we introduce the notion of the minimum coding length of certain asymptotic sufficient statistics and show that it is $\frac{d}{2}\log n + O(1)$, where d is the dimension of the parameter of the family of distribution. Hence, the minimum coding rate is the pre-log coefficient $\frac{d}{2}$, improving over the original d when exact sufficient statistics are used. Here, we also notice that the locality condition can be dropped. That is, our asymptotic sufficient statistics works globally. This is another motivation for the current paper.

A. Related Work

Our problem is different from conventional lossless and lossy source coding [8], [9] because we do not seek to reconstruct the data X^n but rather a distribution on $\mathcal{X}^n$. Hence, we need to generalize standard data compression schemes. Such a generalization has been discussed in the context of quantum data compression by Schumacher [10]. In this setting, the source describing the state cannot be directly observed. Schumacher's encoding process involves compressing the original dataset into a memory with a smaller size. The decoding process involves recovering certain statistics of the data to within a prescribed error bound $\delta \geq 0$.

Recently, Yang, Chiribella and Hayashi [11] extended Schumacher's [10] compression system to a special quantum model. In particular, the authors considered a notion of approximate sufficient statistics in the quantum setting [12], [13] when the data is generated in an i.i.d. manner. They considered only the so-called blind setting [14, Ch. 10] and also only showed a weak converse. We note that there have been recent developments of the notion of approximate sufficient statistics and approximate Markov chains in the quantum information literature [15], [16] but the problem studied here and the objectives are different from the existing works.

Another related line of works in the classical information theory literature are the seminal ones by Rissanen on universal variable-length source coding and model selection [17], [18]. Under the minimum description length (MDL) framework, he introduced a two-step encoding process to obtain a prefix-free source code for n data samples generated from a mixture of i.i.d. distributions. The purpose of Rissanen's compression system is to obtain a compression system for data generated under a mixture distribution. He showed that when the dimensionality of the data is d, the optimal redundancy over the Shannon entropy is $\frac{d}{2}\log n + O(1)$. Merhav and Feder [19] extended Rissanen's analysis to both the minimax and Bayesian (maximin) senses. Clarke and Barron [20], [21] refined Rissanen's analysis and determined the constant (in asymptotic expansions) under various regularity assumptions. While we make heavy use of some of Rissanen's coding ideas and Clarke and Barron's asymptotic expansions for the relative entropy between a parametrized distribution and a mixture, the problem setting we study is different. Indeed, the main ideas in Rissanen's work [17], [18] can only lead to the direct parts with *non-zero* asymptotic error for the relative entropy criterion (see Lemma 1). Similarly, the main results of Clarke and Barron's work [20], [21] can only lead to a *weak converse* under the variational distance criterion (see Lemma 4). Hence, we need to develop new coding techniques and converse ideas to satisfy the more stringent constraints on the code sequences.

B. Main Contributions and Techniques

We provide a precise Shannon-theoretic problem formulation for compression for the model parameter z with an allowable asymptotic error $\delta \geq 0$ on the reconstructed distribution. This error is measured under the relative entropy and variational distance criteria. We use some of Rissanen's ideas for encoding in [17], [18] to show that the memory size can be reduced to approximately $\Theta(n^{\frac{d}{2}})$ resulting in a coding length of $\frac{d}{2}\log n + O(1)$. Note that Rissanen [17], [18] did not explicitly provide the decoders for the problem he considered; we explicitly specify various decoders. Moreover, assuming that the parametric family of distributions is an exponential family [3], we also improve on the evaluations that are inspired by Rissanen (see Lemma 2). In particular, for exponential families, we propose codes whose asymptotic errors measured according to the relative entropy criterion are equal to zero. Furthermore, we consider two separate settings known as the *blind* and *visible* settings. In the former, the encoder can directly observe the dataset X^n; in the latter the encoder directly observes the parameter of interest z. The differences between these two settings are discussed in more detail in [14, Ch. 10]. The visible setting may appear to be less natural but such a generalized setting is useful for the proofs of the converse parts. Yang, Chiribella and Hayashi [11] only considered the blind setting. We consider *both* blind and visible settings and show, somewhat surprisingly, that the coding length is essentially unchanged.

Another significant contribution of our work is in the strengthening of the converse in [11]. In our strong converse proof for the relative entropy error criterion, we employ the Pythagorean theorem for relative entropy, a fundamental concept in information geometry [22]. Furthermore, we use Clarke and Barron's formula [20], [21] to provide a weak converse under the variational distance error criterion. This clarifies the relation between our problem and Clarke and Barron's formula [20], [21]. We significantly strengthen this method to obtain a strong converse (see Lemma 6); in contrast [11] only proves a weak converse. That is, we show that if the error is allowed to be non-vanishing (even if it is arbitrarily large for the relative entropy

criterion and arbitrarily close to 2 for the variational distance criterion), we would still require a memory size of at least $n^{d(\frac{1}{2}-\eta)}$ for any $\eta > 0$ for all sufficiently large n.

C. Paper Organization

This paper is organized as follows. In Section II, we discuss in detail the mathematical formulation of the problem and state the main result as Theorem 1. The results are discussed in the context of exact sufficient statistics and exponential families in Section III. We prove the direct parts of Theorem 1 in Section IV, leveraging ideas from Rissanen's seminal works [17], [18] on universal data compression. We prove the converse parts of Theorem 1 in Section V by leveraging the Pythagorean theorem in information geometry [22] and Clarke and Barron's formula [20], [21].

II. Problem Formulation and Main Results

Let $\mathcal{X}$ be a set and let $\mathcal{P}(\mathcal{X})$ denote the set of distributions (e.g., probability mass functions) on $\mathcal{X}$. We consider a family of distributions $\{P_{X|Z=z}\}_{z\in\mathcal{Z}} \subset \mathcal{P}(\mathcal{X})$ parametrized by a vector parameter $z \in \mathcal{Z} \subset \mathbb{R}^d$. We assume that n independent and identically distributed (i.i.d.) random variables $X^n = (X_1, \ldots, X_n)$ each taking values $\mathcal{X}$ and drawn from $P_{X|Z=z}$. The underlying parameter Z, which is random, follows a distribution $\mu(\mathrm{d}z)$, which is absolutely continuous with respect to the Lebesgue measure on $\mathbb{R}^d$. We will often use the following notations: Given conditional distributions $P_{X|Y}$ and $P_{Y|Z}$, respectively let the joint and marginal probabilities conditioned on $Z = z$ be

$$(P_{X|Y} \times P_{Y|Z})(x, y|z) := P_{X|Y}(x|y)P_{Y|Z}(y|z), \quad \text{and} \tag{2}$$

$$(P_{X|Y} \cdot P_{Y|Z})(x|z) := \sum_y (P_{X|Y} \times P_{Y|Z})(x, y|z). \tag{3}$$

Standard asymptotic notation such as $o(\cdot)$, $O(\cdot)$, and $\Omega(\cdot)$ will be used throughout. Standard information-theoretic notation such as entropy $H(\cdot)$ and mutual information $I(\cdot;\cdot)$ [2] will also be used. Finally, $\|\cdot\|$ and $\|\cdot\|_1$ denote the ℓ_2 and ℓ_1 norms of finite-dimensional vectors respectively.

A. Definitions of Codes

We consider two classes of codes [14, Ch. 10] for the problem of interest:

Definition 1 (Blind code)**.** *A size M_n* blind code *of $\mathcal{C}_{\mathrm{b},n} := (f_{\mathrm{b},n}, \varphi_n)$ consists of*

- *A possibly stochastic encoder (transition kernel) $f_{\mathrm{b},n} : \mathcal{X}^n \to \mathcal{Y}_n := \{1, \ldots, M_n\}$;*
- *A decoder $\varphi_n : \mathcal{Y}_n \to \mathcal{P}(\mathcal{X}^n)$.*

Observe that this definition of a code is similar to that for source coding except that the decoder outputs *distributions* on $\mathcal{X}^n$ instead of length-n strings in $\mathcal{X}^n$. We often consider a more relaxed condition for the encoder as follows. In the visible setting, the encoder does not only have access to the random vector X^n but the parameter $z \in \mathcal{Z}$.

Definition 2 (Visible code)**.** *A size M_n* visible code *$\mathcal{C}_{\mathrm{v},n} := (f_{\mathrm{v},n}, \varphi_n)$ consists of*

- *A possibly stochastic encoder (transition kernel) $f_{\mathrm{v},n} : \mathcal{Z} \to \mathcal{Y}_n := \{1, \ldots, M_n\}$;*
- *A decoder $\varphi_n : \mathcal{Y}_n \to \mathcal{P}(\mathcal{X}^n)$.*

We note that any blind encoder $f_{\mathrm{b},n}$ can be regarded as a special case of a visible encoder $f_{\mathrm{v},n}$ because the visible encoder $f_{\mathrm{v},n}$ can be written in terms of the blind encoder $f_{\mathrm{b},n}$ and the distribution $P^n_{X|Z=z}$ as follows

$$f_{\mathrm{v},n}(z) := \sum_{x^n \in \mathcal{X}^n} f_{\mathrm{b},n}(x^n) P^n_{X|Z=z}(x^n), \quad \forall\, z \in \mathcal{Z}. \tag{4}$$

B. Error Criteria

The performance of any code is characterized by two quantities. First, we desire the coding length $\log M_n = \log |\mathcal{Y}_n|$ to be as short as possible. Next we desire a small error. To define an error criterion precisely, we notice that the reconstructed distribution on $\mathcal{X}^n$ (in the visible case) is $\varphi_n \cdot f_{\mathrm{v},n}(z)$ which is defined as

$$\big(\varphi_n \cdot f_{\mathrm{v},n}(z)\big)(x^n) = \sum_{y \in \mathcal{Y}_n} \Pr\{f_{\mathrm{v},n}(z) = y\} \big(\varphi_n(y)\big)(x^n). \tag{5}$$

This follows from the same notational convention as in (2). Hence the code has a smaller error when the distribution $\varphi_n \cdot f_{\mathrm{v},n}(z)$ is closer to the original distribution $P^n_{X|Z=z}$ for each $z \in \mathcal{Z}$. To evaluate the difference between the two distributions, we

consider an error or fidelity function F whose inputs are distributions on the same probability space. The average error is defined as

$$\varepsilon_{\mathrm{v}}(f_{\mathrm{v},n},\varphi_n) := \int_{\mathcal{Z}} F(\varphi_n \cdot f_{\mathrm{v},n}(z), P^n_{X|Z=z})\,\mu(\mathrm{d}z). \tag{6}$$

For a blind code, in view of (4), we similarly define

$$\varepsilon_{\mathrm{b}}(f_{\mathrm{b},n},\varphi_n) = \varepsilon_{\mathrm{v}}(f_{\mathrm{b},n}\cdot P^n_{X|Z=z},\varphi_n) := \int_{\mathcal{Z}} F(\varphi_n \cdot f_{\mathrm{v},n}\cdot P^n_{X|Z=z}, P^n_{X|Z=z})\,\mu(\mathrm{d}z). \tag{7}$$

It is natural to impose some conditions on the error function F. For this purpose, define the norm $\|z\|_C := \sqrt{\sum_{i,j} z_i z_j C_{ij}}$ for a given positive definite matrix C. We often require F to satisfy

1) $F(P,P') = 0$ if and only if $P = P'$;
2) The approximation
$$F(P^n_{X|Z=z}, P^n_{X|Z=z'}) = n\|z-z'\|_C^2 + o(n\|z-z'\|_C^2) \tag{8}$$
holds for a suitable positive definite matrix C;
3) The joint convexity condition
$$F\left(\sum_i p_i P_i, \sum_i p_i Q_i\right) \le \sum_i p_i F(P_i,Q_i) \tag{9}$$
holds for any probability mass function $\{p_i\}$ and distributions $\{P_i\},\{Q_i\}$.

For example if $F(P,Q)$ is the relative entropy $D(P\|Q) := \sum_x P(x)(\log P(x) - \log Q(x))$, the above conditions are satisfied with C being the Fisher information matrix (see (16) to follow). We also consider the variational distance[1] (also known as the total variation distance) $\|P-Q\|_1 := \sum_x |P(x)-Q(x)|$. Generalizations of these "distances" to continuous-alphabet distributions are performed in the usual manner. We denote the errors in the blind and visible cases [14, Ch. 10] when we use the variational distance as $\varepsilon^{(1)}_{\mathrm{b}}$ and $\varepsilon^{(1)}_{\mathrm{v}}$ respectively. Similarly, we denote the in the blind and visible cases when we use the relative entropy as $\varepsilon^{(2)}_{\mathrm{b}}$ and $\varepsilon^{(2)}_{\mathrm{v}}$ respectively. The size of a code $\mathcal{C}_{\mathrm{b},n}$ is denoted as $|\mathcal{C}_{\mathrm{b},n}| = |\mathcal{Y}_n|$.

C. Definitions of Minimum Compression Rates and Basic Properties

Definition 3 (Minimum Compression Rate)**.** *Let $\delta \ge 0$. We define the* minimum compression rate for blind codes *for a given parametric family $\{P_{X|Z=z}\}_{z\in\mathcal{Z}}$ as*

$$R^{(i)}_{\mathrm{b}}(\delta) := \inf_{\{\mathcal{C}_{\mathrm{b},n}\}_{n\in\mathbb{N}}} \left\{ \varlimsup_{n\to\infty} \frac{\log|\mathcal{C}_{\mathrm{b},n}|}{\log n} \;\middle|\; \varlimsup_{n\to\infty} \varepsilon^{(i)}_{\mathrm{b}}(\mathcal{C}_{\mathrm{b},n}) \le \delta \right\} \tag{10}$$

where $i = 1,2$ denotes whether the error function is the variational distance or relative entropy respectively. In a similar manner, we define the minimum compression rate for visible codes *for a given parametric family $\{P_{X|Z=z}\}_{z\in\mathcal{Z}}$ as*

$$R^{(i)}_{\mathrm{v}}(\delta) := \inf_{\{\mathcal{C}_{\mathrm{v},n}\}_{n\in\mathbb{N}}} \left\{ \varlimsup_{n\to\infty} \frac{\log|\mathcal{C}_{\mathrm{v},n}|}{\log n} \;\middle|\; \varlimsup_{n\to\infty} \varepsilon^{(i)}_{\mathrm{v}}(\mathcal{C}_{\mathrm{v},n}) \le \delta \right\}. \tag{11}$$

To understand this definition, we note that if $R^{(i)}_{\mathrm{b}}(\delta) = c > 0$, then for every $\epsilon > 0$, there exists a sequence of codes $\{\mathcal{C}_{\mathrm{b},n}\}_{n\in\mathbb{N}}$ with asymptotic error no larger than δ and memory or coding length upper bounded as $|\mathcal{C}_{\mathrm{b},n}| \le n^{c+\epsilon}$ for n large enough. Moreover, there is no sequence of codes with with asymptotic error no larger than δ and with $|\mathcal{C}_{\mathrm{b},n}| \le n^{c-\epsilon}$.

The definition of the minimum compression rate differs significantly from traditional source coding in Shannon theory [2] where the normalization of the coding length $\log|\mathcal{C}_{\mathrm{b},n}|$ is n and not $\log n$. Here, we find that the correct normalization is $\log n$. From the above definitions, it is clear that for any $0 \le \delta \le \delta'$ and $i = 1,2$, we have

$$R^{(i)}_a(\delta') \le R^{(i)}_a(\delta), \qquad a = \mathrm{b},\mathrm{v}, \tag{12}$$
$$R^{(1)}_a(0) \le R^{(2)}_a(0), \qquad a = \mathrm{b},\mathrm{v}, \tag{13}$$
$$R^{(i)}_{\mathrm{v}}(\delta) \le R^{(i)}_{\mathrm{b}}(\delta). \tag{14}$$

Note that (13) follows from Pinsker's inequality (i.e., $D(P\|Q) \ge \frac{\log \mathrm{e}}{2}\|P-Q\|_1$) because a vanishing relative entropy implies the same for the variational distance.

[1] We define the variational distance without the coefficient of $\frac{1}{2}$ multiplying $\|P-Q\|_1$ so $0 \le \|P-Q\|_1 \le 2$.

D. Main Results

Let J_z be the Fisher information matrix of the parametric family $\{P_{X|Z=z}\}_{z\in\mathcal{Z}}$. This matrix has elements

$$[J_z]_{i,j} = \mathbb{E}_z\left[\left(\frac{\partial}{\partial z_i}\log P_{X|Z=z}(X)\right)\left(\frac{\partial}{\partial z_j}\log P_{X|Z=z}(X)\right)\right], \tag{15}$$

where $\mathbb{E}_z$ means that we take expectation with respect to $X \sim P_{X|Z=z}$. Before we state the main results of this paper, we consider the following assumptions:

(i) *(Boundedness of Parameter Space)* The set $\mathcal{Z} \subset \mathbb{R}^d$ is bounded;

(ii) *(Euclidean Approximation of Relative Entropy)* As $z' \to z$, the relation

$$D(P_{X|Z=z}\|P_{X|Z=z'}) = \frac{1}{2}\sum_{i,j}[J_z]_{i,j}(z_i - z_i')(z_j - z_j') + o(\|z - z'\|^2) \tag{16}$$

holds [22]–[24]. We also assume compact convergence (i.e., uniform convergence on compact sets) for (16).

(iii) *(Asymptotic Efficiency)* There exists a sequence of estimators $\hat{z}_n = \hat{z}_n(X^n)$ for the parameter z such that

$$\mathbb{E}_z\left[D(P_{X|Z=\hat{z}_n}\|P_{X|Z=z})\right] = \frac{d}{2n} + o\Big(\frac{1}{n}\Big). \tag{17}$$

In other words, the estimator $\hat{z}_n$ asymptotically achieves the Cramér-Rao lower bound [1], [25] (i.e., $\mathbb{E}_z[(\hat{z}_n - z)(\hat{z}_n - z)^T] \to J_z^{-1}$), so the expectation of (16) with $z = z_n$ and $z' = z$ yields (17).

(iv) *(Local Asymptotic Normality)* Fix a point $z \in \mathcal{Z}$ and let $\hat{z}_{\mathrm{ML}}(X^n)$ be the MLE of z given X^n. Define the function $h_z(X^n) = \sqrt{n}J_z^{1/2}(\hat{z}_{\mathrm{ML}}(X^n) - z)$ and let $\phi^{(d)}(x) := (2\pi)^{-d/2}\exp(-\|x\|^2/2)$ be the d-dimensional standard Gaussian probability density function. The *local asymptotic normality* condition [5]–[7] reads

$$\left\|\phi^{(d)} - \left(P^n_{X|Z=z+\frac{z'}{\sqrt{n}}} \cdot h^{-1}_{z+\frac{z'}{\sqrt{n}}}\right)\right\|_1 \to 0 \tag{18}$$

for any vector $z' \in \mathbb{R}^d$.

(v) *(Local Asymptotic Sufficiency)* Let Z be the random variable corresponding to the parameter z and let Y' be the corresponding MLE $\hat{z}_{\mathrm{ML}}(X^n)$. The *local asymptotic sufficiency* condition [5]–[7] reads

$$\left\|\left(P_{X^n|Y',Z=z} \cdot P_{Y'|Z=z+\frac{z'}{\sqrt{n}}}\right) - P^n_{X|Z=z+\frac{z'}{\sqrt{n}}}\right\|_1 \to 0 \tag{19}$$

for any vector $z' \in \mathbb{R}^d$.

Theorem 1. *Assuming (i), (ii), (iv) and (v), the minimum compression rate for visible codes under the variational distance error criterion*

$$R_{\mathrm{v}}^{(1)}(\delta_1) = \frac{d}{2}, \quad \forall\, \delta_1 \in [0, 2). \tag{20}$$

Assuming (i), (ii), the minimum compression rate for visible codes under the relative entropy error criterion

$$R_{\mathrm{v}}^{(2)}(\delta_2) = \frac{d}{2}, \quad \forall\, \delta_2 \in [0, \infty). \tag{21}$$

Assuming (i), (ii), (iv), and (v), the minimum compression rate for blind codes under the variational distance error criterion

$$R_{\mathrm{b}}^{(1)}(\delta_1') = \frac{d}{2}, \quad \forall\, \delta_1' \in [0, 2). \tag{22}$$

Assuming (i), (ii), and (iii) the minimum compression rate for blind codes under the relative entropy error criterion

$$R_{\mathrm{b}}^{(2)}(\delta_2') = \frac{d}{2}, \quad \forall\, \delta_2' \in \Big[\frac{d}{2}, \infty\Big). \tag{23}$$

Furthermore, if (i) holds and $\{P_{X|Z=z}\}_{z\in\mathcal{Z}}$ is an exponential family [3], (23) can be strengthened to

$$R_{\mathrm{b}}^{(2)}(\delta_2') = \frac{d}{2}, \quad \forall\, \delta_2' \in [0, \infty). \tag{24}$$

The direct and converse parts of this theorem are proved in Sections IV and V respectively. Remarks on and implications of the theorem are detailed in the following section.

III. Connection to Sufficient Statistics and Exponential Families

In this section, we discuss the implications of Theorem 1 in greater detail by relating them to the notion of (exact) sufficient statistics [2, Sec. 2.9]. We first review the fundamentals of sufficient statistics, then motivate the notion of approximate sufficient statistics, provide some background on exponential families, and finally show that if one stores the exact sufficient statistics in the memory $\mathcal{Y}_n$, the memory size would be larger than that prescribed by Theorem 1.

A. Review of Sufficient Statistics

Suppose, for the moment, that the blind encoder $f_{\mathrm{b},n}$ is a deterministic function. When $Y = f_{\mathrm{b},n}(X^n)$ is a *sufficient statistic* relative to the family $\{P_{X|Z=z}\}_{z\in\mathcal{Z}}$ [2, Sec. 2.9], the conditional distribution $P^n_{X|Z=z,Y=y}(x^n)$ does not depend on z, i.e., $Z - Y - X$ forms a Markov chain in this order. In this case, we can choose the decoder $\varphi_n : \mathcal{Y}_n \to \mathcal{P}(\mathcal{X}^n)$ as follows

$$\varphi_n(y) := P^n_{X|Z=z,Y=y}. \tag{25}$$

Now, noting that $P^n_{X|Z=z}(\{x^n : f_{\mathrm{b},n}(x^n) = y\}) = f_{\mathrm{b},n} \cdot P^n_{X|Z=z}(y)$ for every y in the memory $\mathcal{Y}_n$, we have

$$\begin{aligned}
\varphi_n \cdot f_{\mathrm{b},n} &= \sum_y f_{\mathrm{b},n} \cdot P^n_{X|Z=z}(y)\varphi_n(y) && (26)\\
&= \sum_y P^n_{X|Z=z}(\{x^n : f_{\mathrm{b},n}(x^n) = y\})P^n_{X|Z=z,Y=y} && (27)\\
&= P^n_{X|Z=z}. && (28)
\end{aligned}$$

Observe that, in this case, regardless of which error metric we choose, we will attain zero error between the reconstructed distribution $\varphi_n \cdot f_{\mathrm{b},n}$ and the original one $P^n_{X|Z=z}$. However, as we will see in a moment, if $P_{X|Z=z}$ is an exponential family the memory size exceeds that prescribed by the various statements in Theorem 1. Thus it is natural to relax the stringent condition in (25) to some approximate versions of it.

B. Exact vs. Approximate Sufficient Statistics

To consider an approximate version of (25), we will consider an additional assumption on the error function F (in addition to the three conditions in Section II-B.

4) Consider distributions P_i and P'_i such that they respectively have disjoint supports from $P_j, j \neq i$ and $P'_j, j \neq i$. Then we assume that

$$F\left(\sum_i p_i P_i, \sum_i p_i P'_i\right) = \sum_i p_i F(P_i, P'_i) \tag{29}$$

where $\{p_i\}$ are non-negative probability masses that sum to one.

Then for $\delta \geq 0$, we consider

$$\begin{aligned}
\varepsilon_{\mathrm{b}}(f_{\mathrm{b},n}, \varphi_n) &= \int_{\mathcal{Z}} F(\varphi_n \cdot f_{\mathrm{b},n} \cdot P^n_{X|Z=z}, P^n_{X|Z=z})\, \mu(\mathrm{d}z) && (30)\\
&= \int_{\mathcal{Z}} F\left(\sum_y (f_{\mathrm{b},n} \cdot P^n_{X|Z=z})(y)\varphi_n(y), \sum_y (f_{\mathrm{b},n} \cdot P^n_{X|Z=z})(y)P^n_{X|Z=z,Y=y}\right) \mu(\mathrm{d}z) && (31)\\
&= \int_{\mathcal{Z}} \sum_y (f_{\mathrm{b},n} \cdot P^n_{X|Z=z})(y) F\big(\varphi_n(y), P^n_{X|Z=z,Y=y}\big)\, \mu(\mathrm{d}z) && (32)\\
&\leq \delta, && (33)
\end{aligned}$$

where (32) uses Condition 4) of F in (29), and (33) uses the fact that the error is bounded by δ according to (10) and (11). If $\delta = 0$, the equality in (25) holds, and we revert to the usual notion of exact sufficient statistics discussed in Section III-A. Hence, the codes that we consider allowing for errors can be regarded as a generalization of sufficient statistics.

C. Background on Exponential Families

To put our results in Theorem 1 into context, we regard $\{P_{X|Z=z}\}_{z\in\mathcal{Z}}$ as an exponential family [3] with parameter space $\mathcal{Z} \subset \mathbb{R}^d$. Recall that a parametric family of distributions $\{P_{X|Z=z}\}_{z\in\mathcal{Z}}$ is called an *exponential family* if it takes the form

$$P_{X|Z=z}(x) = P_X(x) \exp\left[\sum_{i=1}^m z_i Y_i(x) - A(z)\right], \tag{34}$$

where $A(z)$, the log-partition function (cumulant generating function of the random vector $(Y_1(X), \ldots, Y_m(X))$), is defined as

$$A(z) := \log \sum_x P_X(x) \exp\left[\sum_{i=1}^m z_i Y_i(x)\right]. \tag{35}$$

The functions $Y_i(x)$ are known as the *natural statistics* or *sufficient statistics* of the exponential family. Another fact that we exploit in the sequel is that for any exponential family, there is an alternative parametrization known as the *moment parametrization* [3]. There is a one-to-one correspondence between the natural parameter z and the expectation parameter

$$\eta_i(z) := \frac{\partial A(z)}{\partial z_i} = \mathbb{E}_z[Y_i] = \sum_x P_{X|Z=z}(x) Y_i(x). \tag{36}$$

Hence, in the following to estimate the natural parameter z, we can first estimate the moments $\eta(z) = (\eta_1(z), \ldots, \eta_m(z)) \in \mathcal{H} := \{\eta(z) : z \in \mathcal{Z}\}$ and then use the one-to-one correspondence to obtain z.

D. An Example: k-nomial Distributions

Now as a concrete example, we consider a k-nomial distribution, i.e., the family of discrete distributions that take on $k \in \mathbb{N}$ values. The set of k-nomial distributions forms an exponential family with natural statistics

$$Y_i(x) = \begin{cases} 1 & x = i \\ -1 & x = i+1 \\ 0 & \text{else} \end{cases}, \quad i = 1, \ldots, k-1. \tag{37}$$

Note that there are other parametrizations but we work with (37) in the following. It is known that the vector of natural statistics $Y(x) := (Y_1(x), \ldots, Y_{k-1}(x))$ allows us to recover information about the unknown parameter z [22], [26], i.e., $Y(x)$ is sufficient statistic for the k-nomial distribution.

Given n i.i.d. data samples from $P^n_{X|Z=z}$, the exponential family can be written as

$$P^n_{X|Z=z}(x^n) = P^n_X(x^n) \exp\left[\sum_{i=1}^m Y_i^{(n)}(x^n) z_i - nA(z)\right], \tag{38}$$

where $x^n = (x_1, \ldots, x_n) \in \mathcal{X}^n$ and $Y_i^{(n)}(x^n) := \sum_{j=1}^n Y_i(x_j)$. Thus the vector of sufficient statistics is $Y^{(n)}(x^n) := (Y_1^{(n)}(x^n), \ldots, Y_m^{(n)}(x^n))$. In the k-nomial case, the dimension of the exponential family $m = k-1$. It is easy to see that the total number of possibilities of $Y^{(n)}(x^n)$, i.e., the size of the set $\{Y^{(n)}(x^n) : x^n \in \mathcal{X}^n\}$ is $\binom{n+k-1}{k-1}$. This is also the total number of n-types [4] on an alphabet of size k. In this case, the required memory size is

$$\log |\mathcal{Y}_n| = \log \binom{n+k-1}{k-1} = (k-1) \log n + O(1). \tag{39}$$

Note that $k-1 = d$, the dimension of the parameter space $\mathcal{Z}$. Thus, the pre-log coefficient is d, which is twice as large as what the results of Theorem 1 prescribe if we use approximate sufficient statistics in the sense of (33) instead of exact sufficient statistics discussed in Section III-A. This motivates us to study the fundamental limits of approximate sufficient statistics in the large n limit to reduce the memory size $|\mathcal{Y}_n|$ from $n^{d+o(1)}$ to $n^{\frac{d}{2}+o(1)}$.

IV. Proofs of Direct Parts of Theorem 1

In this section, the direct parts (upper bounds) of Theorem 1 will be proved. For logical reasons, the statements in Theorem 1 will not be proved sequentially. Rather we will present the simplest proofs before proceeding to the proofs for more general statements. First, in Section IV-A, we will prove semi-direct part for the relative entropy criterion in (23). This immediately leads the proof of the direct part for (21). Next, in Section IV-B, we strengthen the direct part for exponential families under the relative entropy criterion in (24). Finally, in Section IV-C, we prove the direct part in the blind setting under the variational distance criterion as in (22).

A. Semi-Direct Part Based On Rissanen's Minimum Description Length (MDL) Encoder

Here we prove the direct parts for (21) and (23) where the error criterion used is the relative entropy. We present a complete achievability proof in the visible setting, i.e., (21). Notice that (21) implies (20). Under the same error criterion, we show a semi-achievability in the blind setting (i.e., (23)) in which the error does not vanish even in the limit of large n.

Lemma 1. *Assuming (i), (ii), we have*

$$R_{\mathrm{v}}^{(2)}(0) \le \frac{d}{2}. \tag{40}$$

In addition, assuming (i), (ii), and (iii),

$$R_{\mathrm{b}}^{(2)}\left(\frac{d}{2}\right) \le \frac{d}{2}. \tag{41}$$

Note that (40) proves that (21) holds because it implies that $R_{\mathrm{v}}^{(2)}(\delta_2) \le \frac{d}{2}$ for all $\delta_2 \in [0,\infty)$. Similarly, (41) proves that (23) holds because it implies that $R_{\mathrm{b}}^{(2)}(\delta_2') \le \frac{d}{2}$ for all $\delta_2' \in [\frac{d}{2},\infty)$. These statements follow immediately from the bound in (12) concerning the monotonicity of $\delta \mapsto R_{\mathrm{v}}^{(2)}(\delta)$ and $\delta \mapsto R_{\mathrm{b}}^{(2)}(\delta)$.

We also note that (41), which follows from Rissanen's ideas [17], [18], is rather weak because the asymptotic error is bounded above by $\frac{d}{2}$ instead 0. We improve on this severe limitation in the subsequent subsections.

Proof of Lemma 1: We first prove (41). Then we describe how to modify the argument slightly to show (40). Fix a lattice span $t > 0$ and consider the subset $\mathcal{Z}_{n,t} := \frac{t}{\sqrt{n}}\mathbb{Z}^d \cap \mathcal{Z} \subset \mathcal{Z}$. Given the MLE $\hat{z}_n := \hat{z}_{\mathrm{ML}}(X^n)$, we consider the closest point

$$z_{n,t}(\hat{z}_n) := \operatorname*{arg\,min}_{z' \in \mathcal{Z}_{n,t}} \sum_{i,j} [J_z]_{i,j}(\hat{z}_{n,i} - z_i')(\hat{z}_{n,j} - z_j'). \tag{42}$$

That is, for this blind encoder, the memory $\mathcal{Y}_n$ is taken to be $\mathcal{Z}_{n,t}$ and the encoder is $f_{\mathrm{b},n}(X^n) := z_{n,t}(\hat{z}_{\mathrm{ML}}(X^n))$, i.e., we first compute the MLE then we approximate it with a point in a finite subset $\mathcal{Z}_{n,t}$ using the formula in (42). The decoder is the map from the parameter $z_{n,t}$ to the distribution $P^n_{X|Z=z_{n,t}}$. The coding length is clearly $\log|\mathcal{Z}_{n,t}| = \frac{d}{2}\log n + O(1)$, where the dependence on t is in the $O(1)$ term. Note that Rissanen [17], [18] essentially proposed the same encoder but he was considering a different problem (of universal source coding). Also Rissanen did not explicitly specify the decoder. We also mention that Merhav and Feder [19] extended Rissanen's analysis to both the minimax and Bayesian (maximin) senses.

Now, for any $r > 0$ and any norm $\|\cdot\|$, we have the inequality $2\langle a, b\rangle \le r\|a\|^2 + \frac{1}{r}\|b\|^2$, a consequence of the fact that $\left\|\sqrt{r}a - \sqrt{\frac{1}{r}}b\right\|^2 \ge 0$. Applying this inequality to the norm $\frac{1}{2}\|\cdot\|_{J_z}$ with $a \equiv \hat{z}_n - z$ and $b \equiv z_{n,t} - \hat{z}_n$, we obtain

$$\begin{aligned}
&\frac{1}{2}\sum_{i,j}[J_z]_{i,j}(z_{n,t,i} - z_i)(z_{n,t,j} - z_j) \\
&\quad \le \frac{1+r}{2}\sum_{i,j}[J_z]_{i,j}(\hat{z}_{n,i} - z_i)(\hat{z}_{n,j} - z_j) \\
&\qquad + \frac{1+\frac{1}{r}}{2}\sum_{i,j}[J_z]_{i,j}(z_{n,t,i} - \hat{z}_{n,i})(z_{n,t,j} - \hat{z}_{n,j}).
\end{aligned} \tag{43}$$

We now estimate the error as follows:

$$\varepsilon_{\mathrm{b}}^{(2)}(f_{\mathrm{b},n}, \varphi_n) := \int_{\mathcal{Z}} D\left(\mathbb{E}_z[P^n_{X|Z=z_{n,t}}] \,\middle\|\, P^n_{X|Z=z}\right) \mu(\mathrm{d}z) \tag{44}$$

$$\le \int_{\mathcal{Z}} \mathbb{E}_z\left[D\left(P^n_{X|Z=z_{n,t}} \,\middle\|\, P^n_{X|Z=z}\right)\right] \mu(\mathrm{d}z) \tag{45}$$

$$= \int_{\mathcal{Z}} n\mathbb{E}_z\left[\frac{1}{2}\sum_{i,j}[J_z]_{i,j}(z_{n,t,i} - z_i)(z_{n,t,j} - z_j) + o(\|z_{n,t} - z\|^2)\right] \mu(\mathrm{d}z) \tag{46}$$

$$\begin{aligned}
&\le \int_{\mathcal{Z}} \mathbb{E}_z\Bigg[\frac{n(1+r)}{2}\sum_{i,j}[J_z]_{i,j}(\hat{z}_{n,i} - z_i)(\hat{z}_{n,j} - z_j) \\
&\qquad + \frac{n(1+\frac{1}{r})}{2}\sum_{i,j}[J_z]_{i,j}(z_{n,t,i} - \hat{z}_{n,i})(z_{n,t,j} - \hat{z}_{n,j}) + o(\|z_{n,t} - z\|^2)\Bigg] \mu(\mathrm{d}z)
\end{aligned} \tag{47}$$

$$\begin{aligned}
&= \frac{1+r}{2}\int_{\mathcal{Z}} d\,\mu(\mathrm{d}z) + o(1) \\
&\qquad + \frac{n(1+\frac{1}{r})}{2}\int_{\mathcal{Z}} \mathbb{E}_z\left[\sum_{i,j}[J_z]_{i,j}(z_{n,t,i} - \hat{z}_{n,i})(z_{n,t,j} - \hat{z}_{n,j}) + o\big(\|z_{n,t} - z\|^2\big)\,\mu(\mathrm{d}z)\right]
\end{aligned} \tag{48}$$

$$\le \frac{1+r}{2}\int_{\mathcal{Z}} d\,\mu(\mathrm{d}z) + o(1) + \frac{1+\frac{1}{r}}{2}\sum_{i\ge j}|[J_z]_{i,j}|t^2 + \int_{\mathcal{Z}} \mathbb{E}_z\left[o\big(n\|z_{n,t} - z\|^2\big)\right] \mu(\mathrm{d}z) \tag{49}$$

$$= \frac{(1+r)d}{2} + \frac{1+\frac{1}{r}}{2}\sum_{i\ge j}|[J_z]_{i,j}|t^2 + o(1). \tag{50}$$

We now justify some of the steps above. In (44), the expectation is over the random $z_{n,t}(\hat{z}_{\mathrm{ML}}(X^n))$ where $X^n \sim P^n_{X|Z=z}$; in (45) we used Jensen's inequality and the convexity of the relative entropy; in (46) we used the Euclidean approximation of the relative entropy in (16) (Assumption (ii)); in (47) we used the inequality in (43); in (48) we used (17) (Assumption (iii)); and in (49) and (50) we used the definition of the lattice $\mathcal{Z}_{n,t}$ resulting in the bound $|z_{n,t,i} - \hat{z}_{n,i}| \le \frac{t}{\sqrt{n}}$ for all i and n.

Now since $n \in \mathbb{N}$ is arbitrary,

$$\varlimsup_{n\to\infty} \varepsilon_{\mathrm{b}}^{(2)}(f_{\mathrm{b},n}, \varphi_n) \le \frac{1+r}{2} d + \frac{1+\frac{1}{r}}{2} \sum_{i \ge j} |[J_z]_{i,j}| t^2 \tag{51}$$

Since $t > 0$ is arbitrary, we may take $t \to 0$ so the second term vanishes. Next, since $r > 0$ is arbitrary, we may take $r \to 0$ so the first term converges to the asymptotic error bound of $\frac{d}{2}$. This proves the upper bound in (41).

We now turn our attention to the visible case, which is considerably simpler. In this case, we can replace the MLE $\hat{z}_n$ by z, since the encoder has direct access to the parameter z. Hence, the first terms in (47)–(51) equal 0 and we obtain (40) as desired. ∎

B. Direct Part For Exponential Families

In the blind setting, the MDL encoder discussed in Section IV-A has a non-vanishing error even in the asymptotic limit. To overcome this problem, we devise a novel method attaining zero error in the asymptotic limit. Since the method is more complicated in the general setting, and requires more assumptions (see Section IV-C), we first assume that the distribution family forms an exponential family, and prove the direct part under the relative entropy criterion.

Lemma 2. *When $\{P_{X|Z=z}\}_{z\in\mathcal{Z}}$ is an exponential family and Assumption (i) holds, we have*

$$R_{\mathrm{b}}^{(2)}(0) \le \frac{d}{2}. \tag{52}$$

Note that this indeed implies (24) since $R_{\mathrm{b}}^{(2)}(\delta_2') \le \frac{d}{2}$ for all $\delta_2' \in [0, \infty)$. This significantly improves over the case where we do not assume that $\{P_{X|Z=z}\}_{z\in\mathcal{Z}}$ is an exponential family in (41) of Lemma 1 since we could only prove that $R_{\mathrm{b}}^{(2)}(\delta_2') \le \frac{d}{2}$ for all $\delta_2' \in [\frac{d}{2}, \infty)$, which is much weaker. In other words, the blind code presented below for exponential families can realize the same error performance (which is asymptotically zero) as the visible code presented at the end of the proof of Lemma 1. We also observe that Assumption (i) holds for the k-nomial example discussed in Section III-D as the moment parameters $\mathbb{E}_z[T_i]$ belong to $[-1, 1]$, which is bounded.

Proof of Lemma 2: First, to describe the encoder, we extract the sufficient statistics from the data, i.e., we calculate

$$\hat{\eta}_i := \frac{Y_i^{(n)}(x^n)}{n} = \frac{1}{n} \sum_{j=1}^{n} Y_i(x_j), \quad \forall\, i = 1, \dots, d. \tag{53}$$

Next we fix a lattice span $t > 0$ and consider the subset of quantized moment parameters $\mathcal{H}_{n,t} := \frac{t}{\sqrt{n}} \mathbb{Z}^d \cap \mathcal{H} \subset \mathcal{H}$ where recall that $\mathcal{H} = \{\eta(z) : z \in \mathcal{Z}\}$ is the set of feasible moment parameters (cf. Section III-C). Given the observed value of $\hat{\eta} = (\hat{\eta}_1, \dots, \hat{\eta}_n)$, we choose the closest point in the lattice to it, i.e., we choose

$$\beta_t(\eta) := \operatorname*{arg\,min}_{\eta' \in \mathcal{H}_{n,t}} \|\eta' - \eta\|. \tag{54}$$

The encoder $f_{\mathrm{b},n}$ is simply the map $x^n \mapsto \beta_t(\hat{\eta})$. For this encoder, the memory size is $|\mathcal{H}_{n,t}|$. Thus the coding length is $\log |\mathcal{H}_{n,t}| = \frac{d}{2} \log n + O(1)$, where the dependence in t is in the $O(1)$ term.

Next, we describe the decoder. Let $Y^{(n)} = (Y_1^{(n)}, \dots, Y_d^{(n)})$ and $\eta = (\eta_1, \dots, \eta_d)$. The decoder φ_n is the map

$$\bar{\eta} \mapsto \frac{1}{|\beta_t^{-1}(\bar{\eta})|} \sum_{\eta \in \beta_t^{-1}(\bar{\eta})} P_{X^n|Y^{(n)}=n\eta}. \tag{55}$$

In other words, given the estimate $\beta_t(\eta) \in \mathcal{H}_{n,t}$, we consider a uniform mixture of all the distributions $P_{X^n|Y^{(n)}=n\eta}$ where η runs over all points in the lattice that map to $\beta_t(\eta)$ under the encoding map in (54).

In the following calculation of the error, we first consider the scalar case in which $d = 1$ for simplicity. At the end of the proof, we show how to extend the ideas to the case where $d > 1$. A few additional notational conventions are needed. For each element $\bar{\eta}$, denote the uniform distribution on the subset $\beta_t^{-1}(\bar{\eta})$ as $U_{\beta_t^{-1}(\bar{\eta})}$. Next, denote the transition kernel (channel) that maps the mean parameter $\bar{\eta}$ to the uniform distribution on the set $\beta_t^{-1}(\bar{\eta})$ (i.e., $U_{\beta_t^{-1}(\bar{\eta})}$) as $U_{\beta_t^{-1}(Y)|Y}$. Denote the distribution of the random variable $Y_{n,1}(X^n)$ when $X^n \sim P_{X|Z=z}^n$ as $P_{Y_{n,1}|Z=z}$. Let the variance of $Y_{n,1}$ under distribution $P_{X|Z=z}$ be V_z. The normalizing linear transformation $y \mapsto \sqrt{n}(y - \mathbb{E}_z[Y_{n,1}])/\sqrt{V_z}$ is denoted as g_z.

Since $Y_{n,1}$ is a sufficient statistic relative to the exponential family $\{P_{X|Z=z}\}_{z\in\mathcal{Z}}$, we know that for any error criterion (in particular the relative entropy criterion),

$$\begin{aligned} &D(\varphi_n \cdot f_{\mathrm{b},n} \cdot P_{X|Z=z}^n \,\|\, P_{X|Z=z}^n) \\ &= D(U_{\beta_t^{-1}(Y)|Y} \cdot P_{Y_{n,1}|Z=z} \,\|\, P_{Y_{n,1}|Z=z}) \qquad (56) \\ &= D\big((U_{\beta_t^{-1}(Y)|Y} \cdot P_{Y_{n,1}|Z=z}) \cdot g_z^{-1} \,\|\, P_{Y_{n,1}|Z=z} \cdot g_z^{-1}\big), \qquad (57) \end{aligned}$$

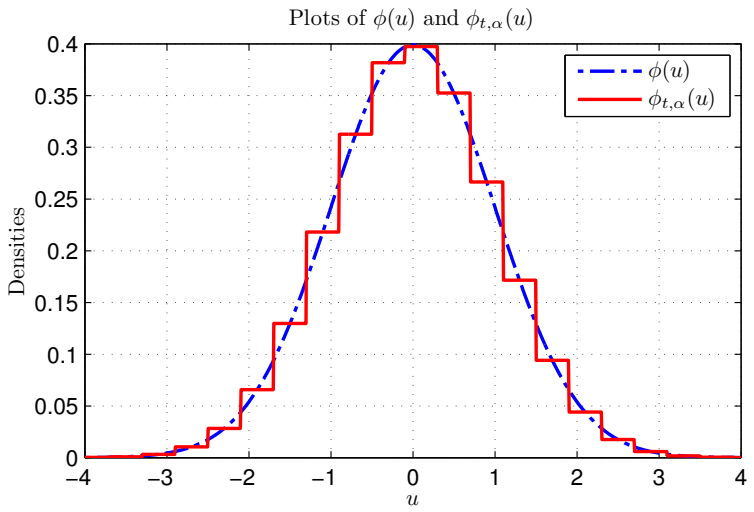

Fig. 1. Illustration of the density $\phi_{t,\alpha}$ in (58) with $t = 0.4$ and $\alpha = 0.3$ (solid red line). The standard normal density ϕ is also shown (broken blue line). The two curves become increasingly close as $t \to 0$.

where the last equality follows from the fact that the function g_z is one-to-one. By the central limit theorem, $P_{Y_{n,1}|Z=z} \cdot g_z^{-1}$ converges to the standard normal distribution $\phi(u) \propto \exp(-u^2/2)$. On the other hand, by the definition of $U_{\beta_t^{-1}(Y)|Y}$ (which results from the construction of the encoder in (54)), the distribution $(U_{\beta_t^{-1}(Y)|Y} \cdot P_{Y_{n,1}|Z=z}) \cdot g_z^{-1}$ converges to a quantization of the standard normal distribution with span t, namely,

$$\phi_{t,\alpha}(u) \propto \phi\big(\alpha + (j+1/2)t\big), \quad \forall\, u \in (\alpha + jt, \alpha + (j+1)t], \tag{58}$$

where $\alpha \in [0, t]$ and the constant of proportionality in (58) is chosen so that $\int_{\mathbb{R}} \phi_{t,\alpha}(u)\,\mathrm{d}u = 1$. See Fig. 1 for an illustration of the probability density function in (58). Thus, we have the upper bound

$$\overline{\lim_{n\to\infty}} D\big((U_{\beta_t^{-1}(Y)|Y} \cdot P_{Y_{n,1}|Z=z}) \cdot g_z^{-1} \,\|\, P_{Y_{n,1}|Z=z} \cdot g_z^{-1}\big) \le \sup_{\alpha\in[0,t]} D(\phi_{t,\alpha} \,\|\, \phi). \tag{59}$$

Since the convergence in (59) is uniform on compact sets (compact convergence), we have

$$\overline{\lim_{n\to\infty}} \varepsilon_{\mathrm{b}}^{(2)}(f_{\mathrm{b},n}, \varphi_n) = \overline{\lim_{n\to\infty}} \int_{\mathcal{Z}} D(\varphi_n \cdot f_{\mathrm{b},n} \cdot P_{X|Z=z}^n \,\|\, P_{X|Z=z}^n)\,\mu(\mathrm{d}z) \tag{60}$$

$$\le \int_{\mathcal{Z}} \overline{\lim_{n\to\infty}} D(\varphi_n \cdot f_{\mathrm{b},n} \cdot P_{X|Z=z}^n \,\|\, P_{X|Z=z}^n)\,\mu(\mathrm{d}z) \tag{61}$$

$$\le \sup_{\alpha\in[0,t]} D(\phi_{t,\alpha} \,\|\, \phi). \tag{62}$$

Now since the above holds for all $t > 0$, we can let t tend to 0 (so the size of the quantization regions decreases to 0). Consequently, the right-hand-side of (62) also tends to 0 and hence, $\lim_{n\to\infty} \varepsilon_{\mathrm{b}}^{(2)}(f_{\mathrm{b},n}, \varphi_n) = 0$.

For general dimension $d > 1$, we can show the desired statement as follows. Let $\phi^{(d)}$ be the d-dimensional standard normal distribution. Given $\alpha \in [0, t]^d$, let $\phi_{t,\alpha}^{(d)}$ be the corresponding quantization of the d-dimensional standard normal distribution with cutting point α and span t (cf. (58) for the one-dimensional distribution). Then in the same way,

$$\overline{\lim_{n\to\infty}} \varepsilon_{\mathrm{b}}^{(2)}(f_{\mathrm{b},n}, \varphi_n) \le \sup_{\alpha\in[0,t]^d} D(\phi_{t,\alpha}^{(d)} \,\|\, \phi^{(d)}). \tag{63}$$

Similarly, we can take t to tend to zero and the error criterion vanishes as $n \to \infty$. The logarithm of the memory size (coding length) is thus $\frac{d}{2}\log n + O(1)$. This proves Lemma 2. ∎

C. Direct Part For The General Case

In this section, we treat the general case (not necessarily exponential family). We prove the following lemma which establishes the direct part (upper bound) in the blind setting under the variational distance criterion as in (22).

Lemma 3. *Assuming (i), (ii), (iv), and (v), we have*

$$R_{\mathrm{b}}^{(1)}(0) \le \frac{d}{2}. \tag{64}$$

Note that this implies the upper bound to (22) because (64) implies that $R_{\mathrm{b}}^{(1)}(\delta_1') \le \frac{d}{2}$ for all $\delta_1' \in [0,2)$.

Proof of Lemma 3: Fix a lattice span $t > 0$ and choose the memory $\mathcal{Y}_n$ to be the quantized parameter space (lattice) $\mathcal{Z}_{n,t} := \frac{t}{\sqrt{n}}\mathbb{Z}^d \cap \mathcal{Z} \subset \mathcal{Z}$. Given the observed MLE $\hat{z}_{\mathrm{ML}}(X^n) = z$, we choose the encoder output to be the closest point in this lattice, i.e.,

$$\beta_t(z) := \operatorname*{arg\,min}_{z' \in \mathcal{Z}_{n,t}} \|z' - z\|. \tag{65}$$

The encoder $f_{\mathrm{b},n}$ is the map from $x^n \mapsto \beta_t(\hat{z}_{\mathrm{ML}}(x^n))$. Thus the code has memory $\mathcal{Y}_n = \mathcal{Z}_{n,t}$ in which the coding length is $\log|\mathcal{Y}_n| = \log|\mathcal{Z}_{n,t}| = \frac{d}{2}\log n + O(1)$.

Now, to describe the decoder and the subsequent analysis, we use some simplified notation. Let Y and Y' denote the random variables $\beta_t(\hat{z}_{\mathrm{ML}}(X^n))$ and $\hat{z}_{\mathrm{ML}}(X^n)$ respectively. These can be thought of as the quantized MLE and the MLE respectively. As usual $Z \in \mathcal{Z}$ is the original parameter. The decoder φ_n is then the following map from elements in the memory to distributions in $\mathcal{P}(\mathcal{X}^n)$:

$$\bar{z} \mapsto \frac{1}{|\beta_t^{-1}(\bar{z})|} \sum_{z \in \beta_t^{-1}(\bar{z})} P_{X^n|Y'=z,Y=\bar{z}}. \tag{66}$$

Essentially, the decoder takes the quantized MLE $\beta_t(\hat{z}_{\mathrm{ML}}(X^n))$ and outputs a uniform mixture over all "compatible" conditional distributions (i.e., all conditional distributions $P_{X^n|Y'=z,Y=\bar{z}}$ whose parameter z lies in the quantization cell $\beta_t^{-1}(\bar{z})$).

In the following, we estimate the error. We first consider the case $d = 1$ for simplicity. For each element $\bar{z}$, denote the uniform distribution on the subset $\beta_t^{-1}(\bar{z})$ as $U_{Y'|Y=\bar{z}}$. We denote the transition kernel corresponding to the map $\bar{z} \mapsto U_{Y'|Y=\bar{z}}$ as $P_{X^n|Y',Y}$. Then the decoder φ_n can alternatively be written as the map $z \mapsto P_{X^n|Y',Y=z} \cdot U_{Y'|Y=z}$. Now the error measured according to the variational distance can be written as

$$\begin{aligned}
\varepsilon_{\mathrm{b}}^{(1)}(f_{\mathrm{b},n}, \varphi_n) &= \int_{\mathcal{Z}} \left\| \varphi_n \cdot f_{\mathrm{b},n} \cdot P_{X^n|Z=z} - P_{X^n|Z=z} \right\|_1 \mu(\mathrm{d}z) && (67)\\
&= \int_{\mathcal{Z}} \left\| (P_{X^n|Y',Y} \cdot U_{Y'|Y}) \times P_{Y=\beta_t(\hat{z}_{\mathrm{ML}}(X^n))|Z=z} - P_{X^n|Z=z} \right\|_1 \mu(\mathrm{d}z) && (68)
\end{aligned}$$

where (68) follows from the definitions of the encoder $P_{Y=\beta_t(\hat{z}_{\mathrm{ML}}(X^n))|Z=z}$ (given the parameter is $Z = z$) and decoder $P_{X^n|Y',Y=z} \cdot U_{Y'|Y=z}$. For clarity, we write $P_{Y=\beta_t(\hat{z}_{\mathrm{ML}}(X^n))|Z=z}$ for the distribution of the quantized MLE $\beta_t(\hat{z}_{\mathrm{ML}}(X^n))$ given that the samples X^n are independently generated from the distribution parametrized by $z \in \mathcal{Z}$, i.e., $P_{X|Z=z}^n$. Let the integrand in (68) for fixed z be denoted as $\varepsilon_{\mathrm{b}}^{(1)}(f_{\mathrm{b},n}, \varphi_n; z)$. By the triangle inequality,

$$\varepsilon_{\mathrm{b}}^{(1)}(f_{\mathrm{b},n}, \varphi_n; z) \le A_n + B_n, \tag{69}$$

where the sequences A_n and B_n are defined as

$$\begin{aligned}
A_n &:= \Big\| (P_{X^n|Y',Y} \cdot U_{Y'|Y}) \times P_{Y=\beta_t(\hat{z}_{\mathrm{ML}}(X^n))|Z=z} \\
&\qquad - P_{X^n|Y',Y} \cdot P_{Y=\beta_t(\hat{z}_{\mathrm{ML}}(X^n)),Y'=\hat{z}_{\mathrm{ML}}(X^n)|Z=z} \Big\|_1, && (70)\\
B_n &:= \left\| P_{X^n|Y',Y} \cdot P_{Y=\beta_t(\hat{z}_{\mathrm{ML}}(X^n)),Y'=\hat{z}_{\mathrm{ML}}(X^n)|Z=z} - P_{X^n|Z=z} \right\|_1. && (71)
\end{aligned}$$

Note that $P_{Y=\beta_t(\hat{z}_{\mathrm{ML}}(X^n)),Y'=\hat{z}_{\mathrm{ML}}(X^n)|Z=z}$ is the joint distribution of the quantized MLE $\beta_t(\hat{z}_{\mathrm{ML}}(X^n))$ and the true MLE $\hat{z}_{\mathrm{ML}}(X^n)$ given that the samples X^n are independently generated from the distribution parametrized by $z \in \mathcal{Z}$. Now, by the data processing inequality for the variational distance (i.e., $\|P_{X|Y} \cdot P_{Y|Z=z} - P_{X|Y} \cdot Q_{Y|Z=z}\|_1 \le \|P_{Y|Z=z} - Q_{Y|Z=z}\|_1$), the term A_n in (70) can be bounded as

$$A_n \le \left\| U_{Y'|Y} \times P_{Y=\beta_t(\hat{z}_{\mathrm{ML}}(X^n))|Z=z} - P_{Y=\beta_t(\hat{z}_{\mathrm{ML}}(X^n)),Y'=\hat{z}_{\mathrm{ML}}(X^n)|Z=z} \right\|_1 \tag{72}$$

We now analyze the right-hand-sides of (72) and (71) in turn.

By a similar reasoning as in the proof of Lemma 2, $U_{Y'|Y} \times P_{Y=\beta_t(\hat{z}_{\mathrm{ML}}(X^n))|Z=z}$ converges to the quantization of the standard normal distribution $\phi_{t,\alpha}$ (defined in (58)) by the local asymptotic normality assumption as stated in (18) (Assumption (iv)). Furthermore, since Y is a deterministic function of Y', we have the relation $P_{Y=\beta_t(\hat{z}_{\mathrm{ML}}(X^n)),Y'=\hat{z}_{\mathrm{ML}}(X^n)|Z=z}(y,y') = P_{Y'=\hat{z}_{\mathrm{ML}}(X^n)|Z=z}(y)\mathbb{1}\{y = \beta_t(y')\}$. Hence, the distribution $P_{Y=\beta_t(\hat{z}_{\mathrm{ML}}(X^n)),Y'=\hat{z}_{\mathrm{ML}}(X^n)|Z=z}$ converges to the standard normal distribution ϕ again by the local asymptotic normality assumption as stated in (18) (Assumption (iv)). Applying the triangle inequality to the right-hand-side of (72), we have

$$\begin{aligned}
\varlimsup_{n\to\infty} A_n &\le \varlimsup_{n\to\infty} \Big\{ \left\| U_{Y'|Y} \times P_{Y=\beta_t(\hat{z}_{\mathrm{ML}}(X^n))|Z=z} - \phi_{t,\alpha} \right\|_1 \\
&\qquad + \left\| \phi_{t,\alpha} - \phi \right\|_1 \\
&\qquad + \left\| \phi - P_{Y=\beta_t(\hat{z}_{\mathrm{ML}}(X^n)),Y'=\hat{z}_{\mathrm{ML}}(X^n)|Z=z} \right\|_1 \Big\} && (73)\\
&\le \sup_{\alpha \in [0,t]} \left\| \phi_{t,\alpha} - \phi \right\|_1. && (74)
\end{aligned}$$

Now we analyze the right-hand-side of (71). First, note that $Y' - z = \hat{z}_{\mathrm{ML}}(X^n) - z$ behaves as $\Theta(\frac{1}{\sqrt{n}})$ with probability tending to one by the central limit theorem (local asymptotic normality). Since the quantization level is also of the order $\Theta(\frac{1}{\sqrt{n}})$, the difference $Y - z = \beta_t(\hat{z}_{\mathrm{ML}}(X^n)) - z$ also behaves as $\Theta(\frac{1}{\sqrt{n}})$ with probability tending to one. Hence, by regarding Y as the random variable $\mathbb{1}\{Z = \tilde{z}\}$ for some $\tilde{z} \in \mathbb{R}^d$ that differs from $z \in \mathcal{Z} \subset \mathbb{R}^d$ by $\Theta(\frac{1}{\sqrt{n}})$, we may write

$$B_n \leq \left\| P_{X^n|Y',Z=\tilde{z}} \cdot P_{Y'=\hat{z}_{\mathrm{ML}}(X^n)|Z=z} - P_{X^n|Z=z} \right\|_1. \tag{75}$$

At this point, we may apply the local asymptotic sufficiency assumption as stated in (19) (Assumption (v)) to (75), yielding

$$\lim_{n\to\infty} B_n = 0. \tag{76}$$

By (69), and similar compact convergence arguments as those leading from (60) to (62), we find that

$$\overline{\lim_{n\to\infty}} \, \varepsilon_{\mathrm{b}}^{(1)}(f_{\mathrm{b},n}, \varphi_n) \leq \sup_{\alpha\in[0,t]} \left\| \phi_{t,\alpha} - \phi \right\|_1. \tag{77}$$

Since this statement holds for all $t > 0$, taking the limit $t \to 0$, we see that the asymptotic error $\lim_{n\to\infty} \varepsilon_{\mathrm{b}}^{(1)}(f_{\mathrm{b},n}, \varphi_n) = 0$. So in the general case for $d = 1$, we can achieve a memory length of $\frac{1}{2}\log n + O(1)$.

The case in which $d > 1$ can be analyzed in a completely analogous manner and we can conclude that a memory length of $\frac{d}{2}\log n + O(1)$ can be achieved. ∎

V. Proofs of Converse Parts of Theorem 1

In this section, we prove the converse parts (lower bounds) to Theorem 1. We will only focus on the visible cases in (20) and (21) because according to (14), a converse for the visible case implies the same for the blind case. Essentially, by (7), a visible code cannot be outperformed by a blind code.

Since our problem seems to be closely related to Clarke and Barron's formula for the relative entropy between a parametrized distribution and a mixture distribution [20], [21], we carefully clarify the relation between our problem and this formula. To clarify this relation, in Section V-A, we prove a weak converse, namely, the impossibility of further compression from a rate of $\frac{d}{2}$ when the variational distance error criterion is asymptotically zero. This can be shown by a simple combination of Clarke and Barron's formula and the uniform continuity of mutual information [27] (also called Fannes inequality [28] in quantum information). Since the variational distance goes to zero when the relative entropy goes to zero, the weak converse under the variational distance criterion implies the weak converse under the relative entropy criterion. Hence, the arguments in Section V-A demonstrate the weak converse under both error criteria. These arguments clarify the relation between Clarke and Barron's formula and our problem. However, to the best of our knowledge, the strong converse parts cannot be shown via Clarke and Barron's formula, i.e., they require novel methods. Furthermore, there is no similar relation between the strong converse parts under the variational distance and the relative entropy. This is because there is no relation between code rates when the relative entropy is arbitrarily large and when the variational distance is arbitrarily close to 2, i.e., its maximum value. So, we need to prove two types of strong converse parts for each of the two error criteria. In Section V-B, we prove a strong converse for the relative entropy error criterion using the Pythagorean theorem for the relative entropy, thus demonstrating (21). In Section V-C, we prove a strong converse for the variational distance error criterion by a different, and novel, method, thus demonstrating (20).

A. Weak Converse Under Both Criteria Based On Clarke And Barron's Formula

In this section, we prove the following weak converse.

Lemma 4. *The following lower bound holds*

$$R_{\mathrm{v}}^{(1)}(0) \geq \frac{d}{2}. \tag{78}$$

This is, in fact, only a weak converse since asymptotically the error measured according to the variational distance must tend to zero. It is insufficient to show (20) but we present the proof to demonstrate the connection between Clarke and Barron's result in (79) to follow and the problem we study. Here, we are only concerned with the variational distance criterion because a weak converse for this criterion implies the same for the relative entropy criterion.

Proof of Lemma 4: We first assume that $\mathcal{X}$ is a finite set. At the end, we show how to relax this condition. We recall that Clarke and Barron [20], [21] showed for a parametric family $\{P_{X|Z=z}\}_{z\in\mathcal{Z}}$ that

$$\begin{aligned} &\int_{\mathcal{Z}} D\left(P_{X|Z=z}^n \,\middle\|\, \int_{\mathcal{Z}} P_{X|Z=z'}^n \, \nu(\mathrm{d}z) \right) \mu(\mathrm{d}z) \\ &= \frac{d}{2}\log\frac{n}{2\pi\mathrm{e}} + D(\mu\|\nu) - D(\mu\|\mu_{\mathrm{J}}) + \log C_{\mathrm{J}} + o(1), \end{aligned} \tag{79}$$

where J_z is the Fisher information matrix defined in (15), $\mu_{\mathrm{J}}(\mathrm{d}z) := \frac{1}{C_{\mathrm{J}}}\det\sqrt{J_z}\,\mathrm{d}z$ is the so-called *Jeffrey's prior* [21] and $C_{\mathrm{J}} := \int_{\mathcal{Z}} \det\sqrt{J_z}\,\mathrm{d}z$ is the normalization factor. When $\nu = \mu$, the left-hand-side of (79) is precisely the mutual information $I(X^n; Z)$ where the pair of random variables (X^n, Z) is distributed according to $P_{X^n,Z}(x^n, z) := P^n_{X|Z=z}(x^n)\mu(z)$. See [29] for an overview of approximations similar to (79) in the context of universal source coding and model order selection.

For the purpose of proving the weak converse, we assume that we are given a sequence of codes $\{\mathcal{C}_{\mathrm{v},n} := (f_{\mathrm{v},n}, \varphi_n)\}_{n\in\mathbb{N}}$ satisfying the condition that the error measured according to the variational distance vanishes, i.e.,

$$\delta_n := \varepsilon^{(1)}_{\mathrm{v}}(\mathcal{C}_{\mathrm{v},n}) \to 0, \quad \text{as } n \to \infty. \tag{80}$$

Now define the code distribution $P_{\mathcal{C}_{\mathrm{v},n}}(x^n, z) := (\varphi_n \cdot f_{\mathrm{v},n}(z))(x^n)\mu(z)$ where $\varphi_n \cdot f_{\mathrm{v},n}(z)$ is defined in (5) and $(\varphi_n \cdot f_{\mathrm{v},n}(z))(x^n)$ is the evaluation of $\varphi_n \cdot f_{\mathrm{v},n}(z)$ at x^n. Then, (80) implies that the variational distance between the code distribution and the generating distribution satisfies

$$\left\| P_{\mathcal{C}_{\mathrm{v},n}} - P_{X^n,Z} \right\|_1 \le \delta_n. \tag{81}$$

Since $|\mathcal{X}|$ is finite, we can use the method of types to find a set of sufficient statistics for the data (cf. Section III-D). Indeed, we can form a set of sufficient statistics relative to the family $\{P_{X|Z=z}\}_{z\in\mathcal{Z}}$. Let us call the sufficient statistics $G_n : \mathcal{X}^n \to \mathbb{R}^{|\mathcal{X}|}$. The output cardinality of G_n is $|G_n(\mathcal{X}^n)| = |\{G_n(x^n) : x^n \in \mathcal{X}^n\}| \le (n+1)^{|\mathcal{X}|-1}$ because we can take the type of x^n to be the sufficient statistic, i.e., $G_n(x^n) = \mathrm{type}(x^n)$. By the data processing inequality for the variational distance, we have

$$\left\| P_{\mathcal{C}_{\mathrm{v},n}} \cdot G_n^{-1} - P_{X^n,Z} \cdot G_n^{-1} \right\|_1 \le \delta_n. \tag{82}$$

In the following we use a subscript to denote the distribution of the random variables in the arguments of the mutual information functional, so for example $I_{P_{AB}}(A; B) = \sum_a P_A(a) D(P_{B|A}(\cdot|a) \| P_B)$. Now, we notice that

$$I_{P_{X^n,Z}\cdot G_n^{-1}}(G_n(X^n); Z) = I_{P_{X^n,Z}}(X^n; Z), \quad \text{and} \tag{83}$$

$$I_{P_{\mathcal{C}_{\mathrm{v},n}}\cdot G_n^{-1}}(G_n(X^n); Z) \le I_{P_{\mathcal{C}_{\mathrm{v},n}}}(X^n; Z) \tag{84}$$

where (83) follows the definition of sufficient statistics and (84) follows from the data processing inequality for mutual information. In addition, by the uniform continuity of mutual information [27], [28] and the estimate in (82), we know that

$$\left| I_{P_{X^n,Z}\cdot G_n^{-1}}(G_n(X^n); Z) - I_{P_{\mathcal{C}_{\mathrm{v},n}}\cdot G_n^{-1}}(G_n(X^n); Z) \right| \le \delta_n \log((n+1)^{|\mathcal{X}|-1}) + \xi(\delta_n), \tag{85}$$

where $\xi(x) := -x\log x$. For brevity, we define the upper bound between the two mutual information quantities in (85) as $\delta'_n := \delta_n \log((n+1)^{|\mathcal{X}|-1}) + \xi(\delta_n)$ and note that $\delta'_n = o(\log n)$.

Next, we define the joint distribution of the encoder and the parameter $P_{f_{\mathrm{v},n}}(y, z) := P_{f_{\mathrm{v},n}(z)}(y)\mu(z)$ (recall the code is visible so the encoder $f_{\mathrm{v},n}$ has access to z). Now we consider the mutual information between the parameter Z and the memory index Y,

$$\begin{aligned}
\log|\mathcal{Y}_n| &\ge I_{P_{f_{\mathrm{v},n}}}(Y; Z) && (86)\\
&= \int_{\mathcal{Z}} D\left(f_{\mathrm{v},n}(z) \,\middle\|\, \int_{\mathcal{Z}} f_{\mathrm{v},n}(z')\,\mu(\mathrm{d}z')\right)\mu(\mathrm{d}z) && (87)\\
&\ge \int_{\mathcal{Z}} D\left(\varphi_n \cdot f_{\mathrm{v},n}(z) \,\middle\|\, \int_{\mathcal{Z}} \varphi_n \cdot f_{\mathrm{v},n}(z')\,\mu(\mathrm{d}z')\right)\mu(\mathrm{d}z) && (88)\\
&= I_{P_{\mathcal{C}_{\mathrm{v},n}}}(X^n; Z) && (89)\\
&\ge I_{P_{\mathcal{C}_{\mathrm{v},n}}\cdot G_n^{-1}}(G_n(X^n); Z) && (90)\\
&\ge I_{P_{X^n,Z}\cdot G_n^{-1}}(G_n(X^n); Z) - \delta'_n && (91)\\
&= I_{P_{X^n,Z}}(X^n; Z) - \delta'_n && (92)\\
&= \int_{\mathcal{Z}} D\left(P^n_{X|Z=z} \,\middle\|\, \int_{\mathcal{Z}} P^n_{X|Z=z'}\,\mu(\mathrm{d}z')\right)\mu(\mathrm{d}z) - \delta'_n && (93)\\
&= \frac{d}{2}\log\frac{n}{2\pi\mathrm{e}} - D(\mu\|\mu_{\mathrm{J}}) + \log C_{\mathrm{J}} + o(1) - \delta'_n && (94)\\
&= \frac{d}{2}\log n + o(\log n). && (95)
\end{aligned}$$

In the above chain, (87), (89), and (93) follow from the definition of mutual information, (88) follows from the data processing inequality for the relative entropy, (90) follows from the data processing inequality for mutual information, (91) follows from the uniform continuity of mutual information as stated in (85), (92) follows from the notion of sufficient statistics as seen in (83), and finally, (94) follows from Clarke and Barron's formula [20] with $\nu = \mu$ in (79). We conclude that if a sequence of codes is such that the variational distance vanishes as in (80), the memory size $|\mathcal{Y}_n|$ must be at least $n^{\frac{d}{2}+o(1)}$.

Now, when $\mathcal{X}$ is not a finite set, we can choose a finite disjoint partition $\{\mathcal{S}_w\}_{w\in\mathcal{W}}$ of $\mathcal{X}$ satisfying the following conditions: (i) $|\mathcal{W}|$ is finite and (ii) $\cup_{w\in\mathcal{W}}\mathcal{S}_w = \mathcal{X}$. Now, we define the parametric family $P_{W|Z=z}(w) := P_{X|Z=z}(\mathcal{S}_w)$. Clearly, we can now go through the above proof with the finite-support random variable W in place of X. Now, when the code reconstructs the original family $\{P^n_{X|Z=z}\}_{z\in\mathcal{Z}}$, clearly it also reconstructs the quantized family $\{P^n_{X|W=w}\}_{w\in\mathcal{W}}$. In essence, reconstructing the latter is "easier" than the former. Since (78) holds for the family $\{P^n_{X|W=w}\}_{w\in\mathcal{W}}$ it must also hold for the family $\{P^n_{X|Z=z}\}_{z\in\mathcal{Z}}$. This completes the proof of (78). ∎

B. Strong Converse Under The Relative Entropy Criterion

In this section, we prove the following strong converse result using the Pythagorean theorem for the relative entropy.

Lemma 5. *The following lower bound holds*

$$R_{\mathrm{v}}^{(2)}(\delta_2) \geq \frac{d}{2}, \quad \forall\, \delta_2 \in [0,\infty). \tag{96}$$

This proves the lower bound to (21). The proof hinges on the Pythagorean formula for the relative entropy and a geometric argument also contained in Rissanen's work [18].

Proof of Lemma 5: Given probability measures $\{P_i\}_{i\in\mathcal{I}}$ and Q, and a probability mass function $\{p_i\}_{i\in\mathcal{I}}$, the Pythagorean formula for relative entropy [22] states that

$$\sum_{i\in\mathcal{I}} p_i D(P_i\|Q) = D\Big(\sum_{i\in\mathcal{I}} p_i P_i\Big\|Q\Big) + \sum_{i\in\mathcal{I}} p_i D\Big(P_i\Big\|\sum_{j\in\mathcal{I}} p_j P_j\Big). \tag{97}$$

In the following, we show that if the memory size $|\mathcal{Y}_n|$ is too small, say $n^{\frac{d}{2}(1-\epsilon)}$ for some fixed $\epsilon > 0$, then the error $\varepsilon_{\mathrm{v}}^{(2)}(\mathcal{C}_{\mathrm{v},n})$ tends to infinity as n grows.

Consider any code $\mathcal{C}_{\mathrm{v},n} = (f_{\mathrm{v},n}, \varphi_n)$ with memory size $|\mathcal{Y}_n| = n^{\frac{d}{2}(1-\epsilon)}$. Let $P_{f_{\mathrm{v},n}(z)}(y) = \Pr\{f_{\mathrm{v},n}(z) = y\}$ be the probability that the index in the memory $Y \in \mathcal{Y}_n$ takes on the value y when the parameter is $z \in \mathcal{Z}$ under the random encoder mapping $f_{\mathrm{v},n}$. Then an application of the Pythagorean theorem in (97) yields

$$\begin{aligned}
&D\Big(\sum_{y\in\mathcal{Y}_n} P_{f_{\mathrm{v},n}(z)}(y)\varphi_n(y)\Big\|P^n_{X|Z=z}\Big)\\
&= \sum_{y\in\mathcal{Y}_n} P_{f_{\mathrm{v},n}(z)}(y) D(\varphi_n(y)\|P^n_{X|Z=z})\\
&\quad - \sum_{y\in\mathcal{Y}_n} P_{f_{\mathrm{v},n}(z)}(y) D\Big(\varphi_n(y)\Big\|\sum_{y'\in\mathcal{Y}_n} P_{f_{\mathrm{v},n}(z)}(y')\varphi_n(y')\Big).
\end{aligned} \tag{98}$$

Hence, by integrating (98) over all $\mathcal{Z}$, we obtain

$$\varepsilon_{\mathrm{v}}^{(2)}(f_{\mathrm{v},n},\varphi_n) = \int_{\mathcal{Z}} D\big(\varphi_n\cdot f_{\mathrm{v},n}(z)\big\|P^n_{X|Z=z}\big)\,\mu(\mathrm{d}z) \tag{99}$$

$$= \int_{\mathcal{Z}} D\Big(\sum_{y\in\mathcal{Y}_n} P_{f_{\mathrm{v},n}(z)}(y)\varphi_n(y)\Big\|P^n_{X|Z=z}\Big)\,\mu(\mathrm{d}z) \tag{100}$$

$$\begin{aligned}
&= \int_{\mathcal{Z}} \sum_{y\in\mathcal{Y}_n} P_{f_{\mathrm{v},n}(z)}(y) D\big(\varphi_n(y)\big\|P^n_{X|Z=z}\big)\,\mu(\mathrm{d}z)\\
&\quad - \int_{\mathcal{Z}} \sum_{y\in\mathcal{Y}_n} P_{f_{\mathrm{v},n}(z)}(y) D\Big(\varphi_n(y)\Big\|\sum_{y'\in\mathcal{Y}_n} P_{f_{\mathrm{v},n}(z)}(y')\varphi_n(y')\Big)\,\mu(\mathrm{d}z).
\end{aligned} \tag{101}$$

We analyze both terms in (101) in turn.

For the first term, we use an argument similar to that for the proof of Theorem 1(a) in Rissanen [18]. Note that since the size of $|\mathcal{Y}_n|$ is $n^{\frac{d}{2}(1-\epsilon)}$, the set $\mathcal{S}_{n,\epsilon}$ of all possible distributions output by the decoder φ_n cannot exceed $n^{\frac{d}{2}(1-\epsilon)}$, i.e., $|\mathcal{S}_{n,\epsilon}| = |\{\varphi_n(y) : y\in\mathcal{Y}_n\}| \leq |\mathcal{Y}_n| = n^{\frac{d}{2}(1-\epsilon)}$. For any given $z\in\mathcal{Z}$, let the closest distribution in $\mathcal{S}_{n,\epsilon}$ have parameter $z'\in\mathcal{Z}$. Since $\mathcal{Z}\in\mathbb{R}^d$ is bounded, we can estimate the (order of the) ℓ_2 distance between z and z', i.e., $\Delta := \|z-z'\|$. If z is a point in general position in $\mathcal{Z}$, then Δ is of the same order as r, where r is the largest radius of the $n^{\frac{d}{2}(1-\epsilon)}$ disjoint spheres contained in $\mathcal{Z}$. Since the volume spheres of radius r in $\mathbb{R}^d$ is proportional to r^d, we have that

$$K_d\cdot r^d\cdot n^{\frac{d}{2}(1-\epsilon)} \geq \mathrm{vol}(\mathcal{Z}), \tag{102}$$

where K_d is a constant that depends only on the dimension d. Since $\mathrm{vol}(\mathcal{Z})$ does not depend on n (it also only depends on d) and $\Delta = \Theta(r)$,[2]

$$\Delta = \Omega(n^{-\frac{1}{2}(1-\epsilon)}). \tag{103}$$

At the same time, by the Euclidean approximation of relative entropy in (16), $D(P^n_{X|Z=z'} \| P^n_{X|Z=z}) = \Omega(n\|z-z'\|^2) = \Omega(n^\epsilon)$. Thus the first term in (101) scales as

$$\int_{\mathcal{Z}} \sum_{y\in\mathcal{Y}_n} P_{f_{\mathrm{v},n}(z)}(y) D\big(\varphi_n(y) \big\| P^n_{X|Z=z}\big)\, \mu(\mathrm{d}z) = \Omega(n^\epsilon). \tag{104}$$

On the other hand the second term in (101) is a conditional mutual information. In particular, it can be upper bounded as

$$\begin{aligned}&\int_{\mathcal{Z}} \sum_{y\in\mathcal{Y}_n} P_{f_{\mathrm{v},n}(z)}(y) D\bigg(\varphi_n(y) \bigg\| \sum_{y'\in\mathcal{Y}_n} P_{f_{\mathrm{v},n}(z)}(y')\varphi_n(y')\bigg)\, \mu(\mathrm{d}z)\\ &= I(X^n; Y|Z) \le H(Y) \le \frac{d}{2}(1-\epsilon)\log n.\end{aligned} \tag{105}$$

Note that the random variables in the information quantities above are computed with respect to the distribution induced by the visible code $\mathcal{C}_{\mathrm{v},n} = (f_{\mathrm{v},n}, \varphi_n)$.

Combining (101), (104), and (105), we obtain

$$\varepsilon^{(2)}_{\mathrm{v}}(f_{\mathrm{v},n}, \varphi_n) \ge \Omega(n^\epsilon) - \frac{d}{2}(1-\epsilon)\log n \to \infty. \tag{106}$$

Hence, with a memory size of $n^{\frac{d}{2}(1-\epsilon)}$, the error computed according to the relative entropy criterion for any visible code diverges. This completes the proof of (96). ∎

C. Strong Converse Under The Variational Distance Criterion

In this section, we prove the following strong converse statement with respect to the variational distance error criterion.

Lemma 6. *The following lower bound holds*

$$R^{(1)}_{\mathrm{v}}(\delta_1) \ge \frac{d}{2}, \quad \forall\, \delta_1 \in [0,2). \tag{107}$$

Lemma 6 significantly strengthens Lemma 4 because the asymptotic error δ_1 is not restricted to be 0; rather it can take any value in $[0,2)$. It demonstrates the lower bound to (20).

Proof of Lemma 6: We first consider the case in which $d = 1$. We proceed by contradiction. We assume, without loss of generality, that the parameter space $\mathcal{Z} = [0,1]$. Fix $\eta \in (0, 1/2)$ and assume that the memory size $M_n = |\mathcal{Y}_n|$ is $O(n^{\frac{1}{2}-\eta})$ and

$$\varepsilon^{(1)}_{\mathrm{v}}(f_{\mathrm{v},n}, \varphi_n) := \mathbb{E}_{z\sim\mu}\left[\big\|\varphi_n \cdot f_{\mathrm{v},n}(z) - P^n_{X|Z=z}\big\|_1\right] \le 2-\alpha \tag{108}$$

for some $\alpha \in (0,2)$ and n large enough. Let

$$\mathcal{S} := \left\{z \in \mathcal{Z} : \big\|\varphi_n \cdot f_{\mathrm{v},n}(z) - P^n_{X|Z=z}\big\|_1 \le 2 - \frac{\alpha}{2}\right\}. \tag{109}$$

Markov's inequality implies that

$$\mu(\mathcal{S}) \ge 1 - \frac{\mathbb{E}_{z\sim\mu}\big[\|\varphi_n \cdot f_{\mathrm{v},n}(z) - P^n_{X|Z=z}\|_1\big]}{2-\frac{\alpha}{2}} \tag{110}$$

$$\ge 1 - \frac{2-\alpha}{2-\frac{\alpha}{2}} \tag{111}$$

$$= \frac{\alpha}{4-\alpha} > 0. \tag{112}$$

Let λ be the Lebesgue measure on $[0,1]$. From (112), we know that $\lambda(\mathcal{S}) > 0$ by absolute continuity of μ with respect to λ (see Section II). Thus, we may choose $\frac{5}{\alpha}M_n$ points $\{z_i : i = 1, \ldots, \frac{5}{\alpha}M_n\} \subset \mathcal{Z}$ satisfying the following two conditions:

$$\big\|\varphi_n \cdot f_{\mathrm{v},n}(z_i) - P^n_{X|Z=z_i}\big\|_1 \le 2 - \frac{\alpha}{2}, \quad \text{and} \tag{113}$$

$$|z_i - z_j| > \lambda(\mathcal{S})\Big(\frac{5}{\alpha}M_n\Big)^{-1}, \quad \forall\, i \ne j. \tag{114}$$

[2]The implied constants in the $\Omega(\,\cdot\,)$ notations used in (103), (104), and (106) are all positive.

Since $\lambda(\mathcal{S})\big(\frac{5}{\alpha}M_n\big)^{-1}$ is $\Omega(n^{-\frac{1}{2}+\eta})$, the distributions $\{P^n_{X|Z=z_i} : i = 1, \ldots, \frac{5}{\alpha}M_n\} \subset \mathcal{P}(\mathcal{X}^n)$ are *distinguishable*. That is, for any $\epsilon > 0$ we may choose an $N \in \mathbb{N}$ satisfying the following. For any $n \geq N$, there exists disjoint subsets $\mathcal{D}_i \subset \mathcal{X}^n$ such that

$$P^n_{X|Z=z_i}(\mathcal{D}_i) \geq 1 - \epsilon \tag{115}$$

for any $i = 1, \ldots, \frac{5}{\alpha}M_n$. For example, we may take

$$\mathcal{D}_i := \left\{ x^n \in \mathcal{X}^n : \Big|\frac{1}{n}\sum_{j=1}^{n} x_j - z_i\Big| \leq \frac{\lambda(\mathcal{S})}{3} \cdot \Big(\frac{5}{\alpha}M_n\Big)^{-1} \right\} \tag{116}$$

and it is then easy to verify that $\{\mathcal{D}_i : i = 1, \ldots, \frac{5}{\alpha}M_n\}$ are disjoint and, by Chebyshev's inequality, that (115) holds for n large enough. Now recall that for any two probability measures P, Q on a common probability space $(\Omega, \mathscr{F})$, half the variational distance can be expressed as $\frac{1}{2}\|P - Q\|_1 = \sup\{P(\mathcal{A}) - Q(\mathcal{A}) : \mathcal{A} \in \mathscr{F}\}$. Thus, the combination of (113) and (115) shows that

$$1 - \frac{\alpha}{4} \geq \big(\varphi_n \cdot f_{\mathrm{v},n}(z_i)\big)(\mathcal{D}_i^c) - P^n_{X|Z=z_i}(\mathcal{D}_i^c) \tag{117}$$

$$\geq \big(\varphi_n \cdot f_{\mathrm{v},n}(z_i)\big)(\mathcal{D}_i^c) - \epsilon. \tag{118}$$

In other words,

$$\big(\varphi_n \cdot f_{\mathrm{v},n}(z_i)\big)(\mathcal{D}_i) \geq \frac{\alpha}{4} - \epsilon. \tag{119}$$

We denote the elements of $\mathcal{Y}_n$ by $\{1, \ldots, M_n\}$. The distribution at the output of the decoder $\varphi_n \cdot f_{\mathrm{v},n}(z)$ is a convex combination of $\{\varphi_n(1), \ldots, \varphi_n(M_n)\}$. Thus,

$$\sum_{j=1}^{M_n} \big(\varphi_n(j)\big)(\mathcal{D}_i) \geq \big(\varphi_n \cdot f_{\mathrm{v},n}(z_i)\big)(\mathcal{D}_i). \tag{120}$$

Hence,

$$M_n \geq \sum_{j=1}^{M_n} \big(\varphi_n(j)\big)\Big(\bigcup_{i=1}^{\frac{5}{\alpha}M_n} \mathcal{D}_i\Big) \tag{121}$$

$$= \sum_{i=1}^{\frac{5}{\alpha}M_n} \sum_{j=1}^{M_n} \big(\varphi_n(j)\big)(\mathcal{D}_i) \tag{122}$$

$$\geq \sum_{i=1}^{\frac{5}{\alpha}M_n} \big(\varphi_n \cdot f_{\mathrm{v},n}(z_i)\big)(\mathcal{D}_i) \tag{123}$$

$$\geq \sum_{i=1}^{\frac{5}{\alpha}M_n} \Big(\frac{\alpha}{4} - \epsilon\Big) \tag{124}$$

$$= \frac{5}{\alpha}M_n\Big(\frac{\alpha}{4} - \epsilon\Big), \tag{125}$$

where (123) and (124) are applications of the bounds in (120) and (119) respectively. So, we obtain

$$1 \geq \frac{5}{\alpha}\Big(\frac{\alpha}{4} - \epsilon\Big), \tag{126}$$

which is a contradiction (if $\epsilon > 0$ is chosen to be smaller than $\frac{\alpha}{20}$). Hence, a memory size of $|\mathcal{Y}_n| = O(n^{\frac{1}{2}-\eta})$ is insufficient to ensure that $\overline{\lim}_{n\to\infty} \varepsilon_{\mathrm{v}}(f_{\mathrm{v},n}, \varphi_n)$ is strictly smaller than 2.

In the general case in which we assume for the sake of contradiction that when the memory size is $|\mathcal{Y}_n| = O(n^{d(\frac{1}{2}-\eta)})$ (for fixed $\eta > 0$), per dimension, the memory size is of the order $O(n^{\frac{1}{2}-\eta})$. Now, we can treat each dimension separately and apply the above argument to yield the same contradiction when the memory size is too small. ■

Acknowledgements

MH is partially supported by a MEXT Grant-in-Aid for Scientific Research (B) No. 16KT0017. MH is also partially supported by the Okawa Research Grant and Kayamori Foundation of Informational Science Advancement. The Centre for Quantum Technologies is funded by the Singapore Ministry of Education and the National Research Foundation as part of the Research Centres of Excellence programme.

VYFT is partially supported an NUS Young Investigator Award (R-263-000-B37-133) and a Singapore Ministry of Education Tier 2 grant "Network Communication with Synchronization Errors: Fundamental Limits and Codes" (R-263-000-B61-112).

REFERENCES

[1] E. L. Lehmann. *Theory of Point Estimation*. Springer, 2nd edition, 1998.
[2] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley-Interscience, 2nd edition, 2006.
[3] M. J. Wainwright and M. I. Jordan. Graphical Models, Exponential Families, and Variational Inference. *Foundations and Trends® in Machine Learning*, 1(1–2):1–305, 2008.
[4] I. Csiszár and J. Körner. *Information Theory: Coding Theorems for Discrete Memoryless Systems*. Cambridge University Press, 2011.
[5] A.W. van der Vaart. *Asymptotic statistics*. Cambridge University Press, 1998.
[6] L. Le Cam. *Asymptotic Methods in Statistical Decision Theory*. Asymptotic Methods in Statistical Decision Theory, 1986.
[7] L. Le Cam. Locally asymptotically normal families of distributions. *University of California Publications in Statistics*, 3:37–98, 1960.
[8] C. E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27:379–423, 1948.
[9] C. E. Shannon. Coding theorems for a discrete source with a fidelity criterion. *IRE Nat. Conv. Rec.*, pages 142–163, 1959.
[10] B. Schumacher. Quantum coding. *Phys. Rev. A*, 51(4):2738–2747, Apr 1995.
[11] Y. Yang, G. Chiribella, and M. Hayashi. Optimal comrpession for identically prepared qubit states. *Phys. Rev. Lett.*, 117(9):090502, Aug 2016.
[12] D. Petz. Sufficient subalgebras and the relative entropy of states of a von Neumann algebra. *Comm. Math. Phys.*, 105(1):123–131, 1986.
[13] M. Koashi and N. Imoto. Compressibility of quantum mixed-state signals. *Phys. Rev. Lett.*, 87(1):017902, 2001.
[14] M. Hayashi. *Quantum Information Theory: A Mathematical Foundation*. Graduate Texts in Physics, Springer, 2nd edition, 2017.
[15] D. Sutter, O. Fawzi, and R. Renner. Universal recovery map for approximate Markov chains. *Proceedings of the Royal Society Series A, Mathematical, Physical and Engineering Sciences*, 472:2186, 2016.
[16] O. Fawzi and R. Renner. Quantum conditional mutual information and approximate Markov chains. *Comm. Math. Phys.*, 340(2):575–611, 2015.
[17] J. Rissanen. A univeral prior for integers and estimation by minimum description length. *Annals of Statistics*, 11(2):416–431, 1983.
[18] J. Rissanen. Universal coding, information, prediction, and estimation. *IEEE Trans. on Inform. Theory*, 30(4):629–636, 1984.
[19] N. Merhav and M. Feder. A strong version of the redundancy-capacity theorem of universal coding. *IEEE Trans. on Inform. Theory*, 41(3):714–722, May 1995.
[20] B. Clarke and A. Barron. Information-theoretic asymptotics of Bayes methods. *IEEE Trans. on Inform. Theory*, 36(3):453–471, 1990.
[21] B. Clarke and A. Barron. Jeffrey's prior is asymptotically least favorable under entropy risk. *J. Statist. Plann. Inference*, 41:37–60, 1994.
[22] S.-I. Amari and H. Nagaoka. *Methods of Information Geometry*. American Mathematical Society, 2000.
[23] S. Borade and L. Zheng. Euclidean Information Theory. In *IEEE International Zurich Seminar on Communications*, pages 14–17, 2008.
[24] E. Abbe and L. Zheng. Linear universal decoding for compound channels. *IEEE Trans. on Inform. Theory*, 56(12):5999–6013, Dec 2010.
[25] I. A. Ibragimov and R. Z. Hasminskii. *Statistical Estimation: Asymptotic Theory*. Springer-Verlag, New York, 1981.
[26] S.-I. Amari. Information geometry on hierarchy of probability distributions. *IEEE Trans. on Inform. Theory*, 47(5):1701–1711, 2001.
[27] Z. Zhang. Estimating mutual information via Kolmogorov distance. *IEEE Trans. on Inform. Theory*, 53(9):3280–82, 2007.
[28] M. Fannes. A continuity property of the entropy density for spin lattice systems. *Comm. Math. Phys.*, 31:291–297, 1973.
[29] A. Barron, J. Rissanen, and B. Yu. The minimum description length principle in coding and modeling. *IEEE Trans. on Inform. Theory*, 44(6):2743–2760, 1998.